

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

1. (Cancelled) A system for allocating a resource to a service request representing a request for a category of service selected from amongst a plurality of possible categories, comprising:

first logic for selecting, responsive to the selected category, a policy from amongst a plurality of possible policies; and

second logic for applying the selected policy to allocate a resource to the request selected from one or more candidate resources.

2. (Cancelled) The system of claim 1 wherein the selected policy is a load balancing policy.

3. (Cancelled) The system of claim 1 wherein the selected category of service is a content-enabled category.

4. (Cancelled) The system of claim 1 wherein the selected category of service is a content-independent category.

5. (Cancelled) The system of claim 1 implemented in hardware as one or more finite state machines.

6. (Cancelled) The system of claim 1 wherein the first logic is configured to determine the policy through an access to a table using an index derived from the service request.

7. (Cancelled) A system for allocating a resource to a service request comprising:

first logic for determining one or more candidate resources using a hierarchical arrangement of data structures, the hierarchical arrangement having a plurality of levels; and

second logic for selecting one of the candidate resources, and allocating the selected resource to the service request.

8. (Cancelled) The system of claim 7 wherein each of the data structures is a table.

9. (Cancelled) The system of claim 8 wherein an index derived from an entry in a table at one level of the hierarchical arrangement is used to obtain an entry in the table at a next level of the hierarchical arrangement.

10. (Cancelled) The system of claim 7 wherein the resource is a server, and the hierarchical arrangement comprises a service index table, a super-group table, and a server group table.

11. (Cancelled) The system of claim 10 wherein an index to the service index table is derived from the service request, and the index is used to access an entry in the service index table specifying a super-group to be allocated to the request, and a load balancing policy.

12. (Cancelled) The system of claim 11 wherein an index derived from the super-group allocated to the request is used to access an entry in the super-group table specifying one or more server groups which are candidates for allocating to the request.

13. (Cancelled) The system of claim 12 wherein one of the candidate server groups is allocated to the request through application of a suitable policy.

14. (Cancelled) The system of claim 13 wherein an index derived from the server group allocated to the request is used to access an entry in the super-group table specifying one or more servers which are candidates for allocating to the request.

15. (Cancelled) The system of claim 14 wherein one of the candidate servers is allocated to the request through application of the load balancing policy specified by the entry in the service index table.

16. (Cancelled) A system for allocating a resource to a service request comprising:

first logic for specifying a plurality of resources which are candidates for allocating to the request; and

second logic for accessing in parallel loading information for each of the candidate resources; and

third logic for allocating one of the candidate resources to the request responsive to the accessed loading information.

17. (Cancelled) The system of claim 16 wherein the loading information for each of the candidate resources is replicated across a plurality of memories, and the second logic is configured to access each of the memories in parallel to obtain the loading information.

18. (Cancelled) The system of claim 16 wherein the third logic is configured to allocate one of the candidate resources to the request through application of a load balancing policy to the loading information for the candidate resources.

19. (Cancelled) A system for allocating a resource to a service request comprising:

first means for determining one or more candidate resources using a hierarchical arrangement of data structures, the hierarchical arrangement having a plurality of levels; and

second means for selecting one of the candidate resources, and allocating the selected resource to the service request.

20. (Cancelled) A method of allocating a resource to a service request representing a request for a category of service selected from amongst a plurality of possible categories, comprising:

determining a policy responsive to the selected category; and

applying the policy to allocate a resource to the request selected from one or more candidate resources.

21. (Cancelled) The method of claim 20 wherein the policy is a load balancing policy.

22. (Cancelled) The method of claim 20 wherein the policy is determined through an access to a table using an index derived from the selected category of service.

23. (Cancelled) A method of allocating a resource to a service request comprising:

determining one or more candidate resources using a hierarchical arrangement of data structures, the hierarchical arrangement having a plurality of levels; and

selecting one of the candidate resources, and allocating the selected resource to the service request.

24. (Cancelled) The method of claim 23 wherein each of the data structures in the hierarchical arrangement is a table.

25. (Cancelled) The method of claim 24 further comprising deriving an index from an entry in a table at one level of the hierarchy, and using the index to access an entry in a table at a next level of the hierarchy.

26. (Cancelled) The method of claim 23 wherein the determining step comprises:

deriving an index to a service index table from the service request;
using the index to access an entry in the service index table; and
allocating a super-group to the request and determining a load balancing policy responsive to the entry in the service index table.

27. (Cancelled) The method of claim 26 wherein the resource is a server, and the determining step further comprises:

deriving an index to a super-group table from the super-group allocated to the request;
using the index to access an entry in the super-group table;
determining from the entry one or more server groups which are candidates for allocating to the request; and
allocating one of the candidate server groups to the request.

28. (Cancelled) The method of claim 27 further comprising allocating one of the candidate server groups to the request through application of a suitable load balancing policy.

29. (Cancelled) The method of claim 27 wherein the determining step further comprises:

deriving an index to a server group table from the server group allocated to the request;

using the index to access an entry in a server group table; and

determining from the entry the one or more servers which are candidates for allocating to the request.

30. (Cancelled) The method of claim 23 wherein the resource is a server, further comprising allocating the selected server to the request only if a server is not allocated to the request through application of a persistence policy.

31. (Cancelled) A method of allocating a resource to a service request comprising:

a step for determining one or more candidate resources using a hierarchical arrangement of data structures, the hierarchical arrangement having a plurality of levels; and

a step for selecting one of the candidate resources, and allocating the selected resource to the service request.

32. (Cancelled) A method of allocating a resource to a service request comprising:

specifying a plurality of resources which are candidates for allocating to the request; and

accessing in parallel loading information for each of the candidate resources; and

allocating one of the candidate resources to the request responsive to the accessed loading information.

33. (Cancelled) The method of claim 32 wherein the resource is a server, and the accessing step comprises accessing in parallel the loading information from a server loading table replicated across a plurality of memories which are accessible in parallel.

34. (Cancelled) The method of claim 32 wherein the allocating step comprises allocating one of the candidate resources to the request responsive to application of a load balancing policy to the accessed loading information.

35. (Cancelled) The system of claim 1 wherein the resource is a server.

36. (Cancelled) The system of any of claims 7 or 19 wherein the one or more candidate resources are servers.

37. (Cancelled) The method of claim 20 wherein the resource is a server.

38. (Cancelled) The method of any of claims 23 or 31 wherein the one or more candidate resources are servers.

39. (Cancelled) The method of claim 32 wherein the plurality of resources are servers.

40. (Cancelled) The system of any of claims 1, 7, 16, or 19, wherein the resources are servers.

41. (Cancelled) The system of any of claims 1, 7, 16, or 19, wherein the service requests are in the form of or spawned by packets.

42. (Cancelled) The system of any of claims 1, 7, 16, or 19, implemented as one or more engines.

43. (Cancelled) The method of any of claims 20, 23, 31, and 32, wherein the resources are servers.

44. (Cancelled) The method of any of claims 20, 23, 31, and 32, wherein the service requests are in the form of or spawned by packets.

45. (New) A system for allocating resources to service requests comprising:

a service index block having a plurality of super group indexes, the service index block being coupled to receive a first service index corresponding to a service request and configured to provide a first super group index in response to the first service index;

a super group block having a plurality of server group indexes, the super group block being coupled to receive the first super group index and configured to provide a first resource group index in response to the first super group index;

a group block having a plurality of resource indexes, the group block being coupled to receive the first resource group index and configured to provide at least one resource index in response to the first resource group index, each resource index corresponding to one of a plurality of resources.

46. (New) The system of Claim 45 wherein the at least one resource index is one resource index, the one resource index indicating a resource assigned to the service request.

47. (New) The system of Claim 46 wherein the resource is a server.

48. (New) The system of Claim 45 wherein the service index block further comprises a plurality of load balancing policy indexes, the service index block further configured to provide a first load balancing policy index in response to the first service index, the system further comprising a first load balancer coupled to receive the at least one resource index and the first load balancing policy index, select a first load balancing policy from among a plurality of load balancing policies based on the first load balancing policy index, and apply the first load balancing policy to the at least one resource index to select a resource.

49. (New) The system of Claim 48 wherein the system further comprises a second load balancer, the super group block further configured to generate at least one additional resource group index in response to the first super group index, the second load balancer configured to select the first resource group index from among the first resource group index and the at least one additional resource group index based on a second load balancing policy.

50. (New) The system of claim 49 wherein the second load balancing policy is selected based on the first load balancing policy index.

51. (New) The system of claim 49 wherein the second load balancing policy is configured independently of the first load balancing policy index.

52. (New) The system of Claim 45 further comprising a content analysis engine configured to receive at least a portion of a data packet and generate a service index based on at least one of domain name and URL pattern matching.

53. (New) The system of Claim 45 further comprising a lookup table coupled to receive at least a portion of a data packet and configured to lookup the first service index based on at least one of a destination IP, a destination port and a protocol corresponding to the data packet.

54. (New) The system of Claim 45 further comprising a history table configured to receive the first service index and select a second resource index based on a persistence policy if a match is found in the history table.

55. (New) The system of Claim 45 wherein the group block comprises a hierarchical data structure using the first resource group index to produce the at least one resource index, the hierarchical data structure comprising a plurality of tables each accessed in sequence, an output of each table being an input to the next table in the sequence, the first sequential table of the plurality of tables being configured to receive the first resource group index and the last sequential table of the plurality of tables being configured to provide the at least one resource index.

56. (New) A method for allocating resources to service requests comprising:

receiving a first service index corresponding to a service request;

generating a first super group index in response to the first service

index;

generating a first resource group index in response to the first super group index; and

generating at least one resource index in response to the first resource group index, the at least one resource index corresponding to at least one of a plurality of resources.

57. (New) The method of Claim 56 further comprising the step of assigning a resource to the service request, wherein the at least one resource index is one resource index, the one resource index indicating the resource.

58. (New) The method of Claim 57 wherein the resource is a server.

59. (New) The method of Claim 57 further comprising:

generating a first load balancing policy index in response to the first service index;

selecting a first load balancing policy from among a plurality of load balancing policies based on the first load balancing policy index; and

applying the first load balancing policy to the at least one resource index to select a resource.

60. (New) The method of Claim 59 further comprising:

generating at least one additional resource group index in response to the first super group index; and

applying a second load balancing policy to the first resource group index and the at least one additional resource group index to select the first resource group index.

61. (New) The method of Claim 60 further comprising selecting the second load balancing policy based on the first load balancing policy index.

62. (New) The method of Claim 60 further comprising selecting the second load balancing policy independently of the first load balancing policy index.

63. (New) The method of Claim 56 further comprising:
receiving at least a portion of a data packet having a domain name and a URL; and
generating a service index based on the domain name and pattern matching of within the URL.

64. (New) The method of Claim 56 further comprising:
receiving at least a portion of a data packet having a destination IP, a destination port and a protocol; and
generating a service index based on at least one of the destination IP, the destination port and the protocol.

65. (New) The method of Claim 56 further comprising selecting a second server based on a persistence policy if a match is found in a history table.

66. (New) The method of Claim 65 wherein the steps of generating the first super group index, generating a first group index and generating a first server being performed only if a match is not found in the history table.

67. (New) The system of Claim 56 wherein generating at least one resource index in response to the first resource group index comprises accessing a plurality of tables in sequence, the first sequential table of the plurality of tables receiving the first resource group index, each of the plurality of tables providing an

input to the next table in the sequence, and the last sequential table of the plurality of tables providing the at least one resource index.